SEPARATING HYPERPLANES OF EDGE POLYTOPES

TAKAYUKI HIBI, NAN LI AND YAN X ZHANG

ABSTRACT. Let G be a finite connected simple graph with d vertices and let $\mathcal{P}_G \subset \mathbb{R}^d$ be the edge polytope of G. We call $\mathcal{P}_G$ *decomposable* if $\mathcal{P}_G$ decomposes into integral polytopes $\mathcal{P}_{G^+}$ and $\mathcal{P}_{G^-}$ via a hyperplane. In this paper, we explore various aspects of decomposition of $\mathcal{P}_G$: we give an algorithm deciding the decomposability of $\mathcal{P}_G$, we prove that $\mathcal{P}_G$ is normal if and only if both $\mathcal{P}_{G^+}$ and $\mathcal{P}_{G^-}$ are normal, and we also study how a condition on the toric ideal of $\mathcal{P}_G$ (namely, the ideal being generated by quadratic binomials) behaves under decomposition.

INTRODUCTION

A *simple graph* is a graph with no loops and no multiple edges. Let G be a finite connected simple graph with vertex set $[d] = \{1, \ldots, d\}$ and edge set $E(G) = \{e_1, \ldots, e_n\}$. Let $\mathbf{e}_i$ be the i-th unit coordinate vector of the euclidean space $\mathbb{R}^d$. If $e = (i, j)$ is an edge of G, then we set $\rho(e) = \mathbf{e}_i + \mathbf{e}_j \in \mathbb{R}^d$. The *edge polytope* $\mathcal{P}_G$ of G is the convex hull of $\{\rho(e_1), \ldots, \rho(e_n)\}$ in $\mathbb{R}^d$. The basics of edge polytopes are studied in [4], [5], [6], [7] and [8].

In this paper, we study the decompositions of edge polytopes via hyperplanes. Recall that a convex polytope is *integral* if all of its vertices have integral coordinates; in particular, $\mathcal{P}_G$ is an integral polytope. Let ∂P denote the boundary of a polytope P. We say that P is *decomposable* if there exists a hyperplane $\mathcal{H}$ of $\mathbb{R}^d$ with $\mathcal{H} \cap (P \setminus \partial P) \neq \emptyset$ such that each of the convex polytopes $P \cap \mathcal{H}^{(+)}$ and $P \cap \mathcal{H}^{(-)}$ is integral. Here $\mathcal{H}^{(+)}$ and $\mathcal{H}^{(-)}$ are the closed half-spaces of $\mathbb{R}^d$ with $\mathcal{H}^{(+)} \cap \mathcal{H}^{(-)} = \mathcal{H}$. Such a hyperplane $\mathcal{H}$ is called a *separating hyperplane* of P.

We start with a nice fact, Lemma 1.8, which shows that when $\mathcal{P}_G$ is decomposable into $\mathcal{P}_{G^+}$ and $\mathcal{P}_{G^-}$, the two subpolytopes are again edge polytopes. This lemma sets the theme for our paper, namely that certain polytope properties (in this case, being an edge polytope) are well-behaved under decomposition when we restrict to edge polytopes. Furthermore, thanks to edge polytopes having associated graphs, these properties may correspond to easily-visualized combinatorial properties of the underlying graphs. When both of these conditions are in place, checking difficult properties on a decomposable polytope may then be reduced to checking those properties on the graphs corresponding to the smaller pieces in its decomposition. This is one motivation for the study of decomposable edge polytopes.

After the basics, we look at the fundamental problem of determining which edge polytopes are decomposable. In Corollary 1.6, we see that a necessary condition for the decomposability of $\mathcal{P}_G$ is that G possesses at least one cycle of length 4. We also provide Algorithm 1.13, which decides the decomposability of an edge polytope. We then focus on the case when G is a complete multipartite graph and count the number of separating hyperplanes for such $\mathcal{P}_G$ in Theorem 2.5.

Then, we consider the property of normality under decomposition. The characterization of normal polytopes is one of the fundamental questions of the study of lattice polytopes that can also shed light on other properties. For example, it is hard to check combinatorially whether a toric ring of a polytope is Cohen-Macaulay, but Hochster's Theorem [3] gives that this is implied by the normality of the polytope. This paper is part of a continuing effort to understand normality in edge polytopes that started in [4] and [10]. It is known [4, Corollary 2.3] that an edge polytope $\mathcal{P}_G$ is normal if and only if G satisfies the so-called "odd cycle condition" ([1]). We use this combinatorial criterion to show that normality of edge polytopes also behaves nicely under decomposition in Theorem 3.2; specifically, in a decomposition of an edge polytope $\mathcal{P}_G$, then $\mathcal{P}_G$ is normal if and only if both subpolytopes are normal.

Finally, we examine the connections between toric ideals and lattice polytopes. In the theory of toric ideals ([11]), special attention has been given to toric ideals generated by quadratic binomials. As before, we study this property under decomposition by checking a combinatorial condition on the graph, in this case with the help of [7]. We show in Theorem 4.3 that I_G is generated by quadratic binomials if both toric ideals corresponding to the subpolytopes have this property. However, we stress that the converse does not hold.

There are other similar questions to be asked following our work. For example, we can ask how the property of the toric ring possessing a quadratic (or squarefree) Gröbner basis behaves under decomposition. Unfortunately, we currently do not have a combinatorial description for this property and thus do not know how to approach this problem.

1. SEPARATING HYPERPLANES AND DECOMPOSITIONS

The vertices of the edge polytope $\mathcal{P}_G$ of G are $\rho(e_1), \ldots, \rho(e_n)$, but not all edges of the form $(\rho(e_i), \rho(e_j))$ actually occur. For $i \neq j$, let $\mathrm{co}(e_i, e_j)$ be the convex hull of the pair $\{\rho(e_i), \rho(e_j)\}$. The edges of $\mathcal{P}_G$ will be a subset of these $\mathrm{co}(e_i, e_j)$. For edges $e = (i, j)$ and $f = (k, \ell)$, call the pair of edges (e, f) *cycle-compatible* with C if there exists a 4-cycle C in the subgraph of G induced by $\{i, j, k, \ell\}$ (in particular, this implies that e and f do not share any vertices). The following result allows us to identify the $\mathrm{co}(e_i, e_j)$ that are actually edges of $\mathcal{P}_G$ using the notion of cycle-compatibility.

Lemma 1.1 ([9]). *Let e and f be edges of G with $e \neq f$. Then $\mathrm{co}(e, f)$ is an edge of $\mathcal{P}_G$ if and only if e and f are not cycle-compatible.*

Example 1.2. Let K_d denote the complete graph on $[d]$ and e and f be edges of K_d. We have that $e \cap f = \emptyset$ exactly when e and f are cycle-compatible. We can then compute the number of edges of $\mathcal{P}_{K_d}$ by counting the 2-element subsets $\{e, f\}$ of $E(K_d)$ with $e \cap f \neq \emptyset$, which is

$$d \binom{d-1}{2} = \frac{d^2(d-1)}{2}.$$

Question 1.3. Fix integers $d \geq 2$. What is the maximal number of possible edges of $\mathcal{P}_G$, as G ranges among all finite connected simple graphs on $[d]$?

The condition that $\mathcal{P}_G \cap \mathcal{H}^{(+)}$ and $\mathcal{P}_G \cap \mathcal{H}^{(-)}$ are integral is equivalent to having no edge in $\mathcal{P}_G$ intersecting $\mathcal{H}$ except possibly at the endpoints. This is, by Lemma 1.1, equivalent to the the following condition: we say that G satisfies the *4-cycle condition* if for any pair of edges $e, f \in E(G)$ such that $\rho(e) \in \mathcal{H}^{(+)} \backslash \mathcal{H}$ and $\rho(f) \in \mathcal{H}^{(-)} \backslash \mathcal{H}$, e and f are cycle-compatible.

We can make a simplification of $\mathcal{H}$:

Lemma 1.4. *If $\mathcal{P}_G$ is decomposable via a hyperplane $\mathcal{H}$ that does not go through the origin, there exists a hyperplane $\mathcal{H}'$ that gives the same decomposition, with the additional condition that $\mathcal{H}$ goes through the origin.*

Proof. Suppose $\mathcal{H} \subset \mathbb{R}^d$ is defined by

$$a_1 x_1 + \cdots + a_d x_d = b,$$

with each $a_i, b \in \mathbb{R}$. Let $\mathcal{H}'$ denote the hyperplane defined by the equation

$$(a_1 - b/2)x_1 + \cdots + (a_d - b/2)x_d = 0.$$

Since $\mathcal{H} \cap \mathcal{P}_G$ lies in the hyperplane defined by the equation $x_1 + \cdots + x_d = 2$, it follows that $\mathcal{H} \cap \mathcal{P}_G = \mathcal{H}' \cap \mathcal{P}_G$ and the decomposition is not affected. $\square$

Lemma 1.4 allows us to assume that all our separating hyperplanes for edge polytopes go through the origin; we will always make this assumption from now on. When we do, assume that $\mathcal{H}^{(+)}$ contains points $(x_1, \ldots, x_n)$ where $\sum a_i x_i \geq 0$ and $\mathcal{H}^{(-)}$ contains points where $\sum a_i x_i \leq 0$.

We now introduce the function $s_{\mathcal{H}} : E(G) \to \{0, 1, -1\}$ defined by setting $s_{\mathcal{H}}((i, j))$ to be the sign of $a_i + a_j$, allowing 0. The function $s_{\mathcal{H}}$ enables us to call an edge e "positive," "negative," or "zero," corresponding to whether the associated vertex $\rho(e)$ in $\mathcal{P}_G$ is in $\mathcal{H}^{(+)} \backslash \mathcal{H}$, $\mathcal{H}^{(-)} \backslash \mathcal{H}$, or $\mathcal{H}$, respectively. We will repeatly use the following fact:

Corollary 1.5. *For any positive edge e and negative edge f in a decomposition, e and f must be cycle-compatible in a cycle (e, g, f, h), where g and h are zero edges.*

Proof. The cycle-compatibility is an immediate corollary of Lemma 1.1. This also implies we cannot have a positive edge sharing a vertex with a negative edge, so since the other two edges share a vertex with both e and f, they must be zero edges. $\square$

Since the hyperplane $\mathcal{H}$ decomposes $\mathcal{P}_G$, we must have at least one positive edge and at least one negative edge. Thus, we also have:

Corollary 1.6. *Suppose that $\mathcal{P}_G$ is decomposable. Then G must possess at least one cycle of length* 4.

Example 1.7. The converse of Corollary 1.6 does not hold for all graphs. Let G be the following graph.
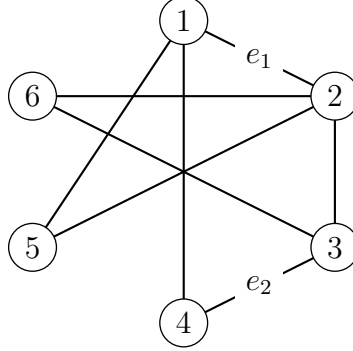
Even though G possesses a cycle of length 4, the edge polytope $\mathcal{P}_G$ is indecomposable. If $\mathcal{H}$ were a separating hyperplane with equation $\sum_i^6 a_i x_i = 0$, since there is only one 4-cycle and the two pairs of nonadjacent edges are symmetric, we may assume without loss of generality that e_1 is positive and e_2 is negative. Since any pair of positive and negative edges must form a 4-cycle and there are no other 4-cycles, all remaining e_i must be zero edges. Therefore, we have

$$a_2 + a_3 = a_1 + a_4 = a_1 + a_5 = a_2 + a_5 = a_2 + a_6 = a_3 + a_6 = 0.$$

Thus $a_1 = a_2 = a_3 = -a_4 = -a_5 = -a_6$. In particular $a_3 + a_4 = 0$. However, since $\rho(e_2) \in \mathcal{H}^{(-)} \setminus \mathcal{H}$, one has $a_3 + a_4 \neq 0$, a contradiction.

The following important result tells us that being an edge polytope is hereditary under decomposition:

Lemma 1.8. *Let G be a finite connected simple graph on $[d]$ and suppose that $\mathcal{P}_G \subset \mathbb{R}^d$ is decomposable by $\mathcal{H}$. Then each of the subpolytopes $\mathcal{P}_G \cap \mathcal{H}^{(+)}$ and $\mathcal{P}_G \cap \mathcal{H}^{(-)}$ is again an edge polytope. More precisely, one has connected spanning subgraphs G_+ and G_- with $\mathcal{P}_G \cap \mathcal{H}^{(+)} = \mathcal{P}_{G_+}$, $\mathcal{P}_G \cap \mathcal{H}^{(-)} = \mathcal{P}_{G_-}$.*

Proof. Let G_+ and G_- be subgraphs of G with $E(G_+) = \{e \in E(G) : s_{\mathcal{H}}(e) \geq 0\}$ and $E(G_-) = \{e \in E(G) : s_{\mathcal{H}}(e) \leq 0\}$. Let $\mathcal{P}_{G_+}$ be the subpolytope of $\mathcal{P}_G$ which is the convex hull of $\{\rho(e) : e \in E(G_+)\}$ and let $\mathcal{P}_{G_-}$ be the subpolytope of $\mathcal{P}_G$ which

is the convex hull of $\{\rho(e) : e \in E(G_-)\}$. One can assume $\mathcal{P}_G \cap \mathcal{H}^{(+)} = \mathcal{P}_{G_+}$ and $\mathcal{P}_G \cap \mathcal{H}^{(-)} = \mathcal{P}_{G_-}$.

Since the dimension of each of the subpolytopes $\mathcal{P}_{G_+}$ and $\mathcal{P}_{G_-}$ coincides with that of $\mathcal{P}_G$, it follows from [4, Proposition 1.3] that both subgraphs G_+ and G_- must be connected spanning subgraphs of G, as desired. $\square$

Example 1.9. An example of a decomposition of a graph G using the hyperplane $-x_1 + x_4 - x_5 + x_6 = 0$. For edges (i, j) in G, $a_i + a_j$ is only nonzero for $\{i, j\} = \{1, 2\}$ or $\{3, 4\}$, with values -1 and 1 respectively. These correspond to the two non-zero edges. We give two equivalent graphical representations of this decomposition, one by showing G_+ and G_- explicitly and one by marking the non-zero edges $+$ or $-$.
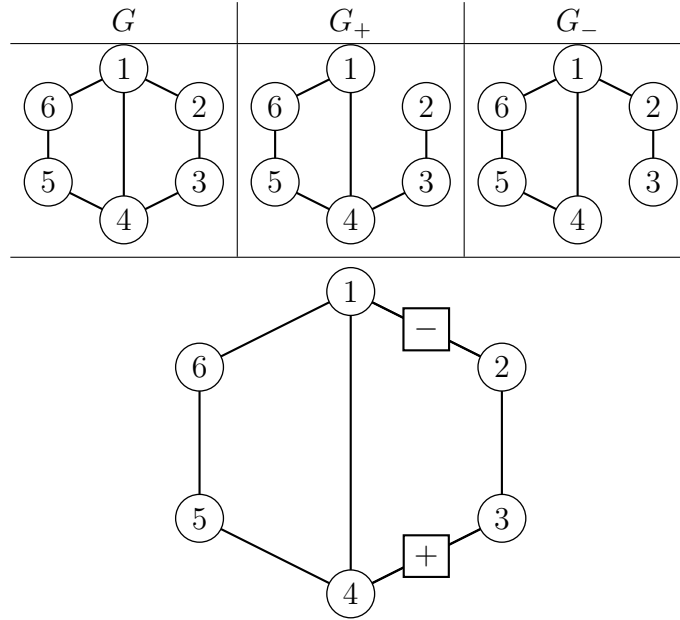
We will now prove a result that allows us to normalize the coefficients used in the hyperplane in a decomposition. Let the *weight* of a vertex i of G be a_i and let the *signature* of the edge (i, j) (or the corresponding vertex $\rho((i, j))$ in $\mathcal{P}_G$) be the set of weights $\{a_i, a_j\}$ and the *weight* of (i, j) be $a_i + a_j$. An edge e has the sign function $s_{\mathcal{H}}(e)$ equal to the sign of the sum of the weights in its signature.

Proposition 1.10. *Suppose we have a decomposition $\mathcal{P}_G = \mathcal{P}_G^{(+)} \cup \mathcal{P}_G^{(-)}$. Then for some fixed a, b with $a + b > 0$, every positive edge must have signature $\{a, b\}$ and every negative edge must have signature $\{-a, -b\}$. Furthermore, we can assume that all weights take values in $\{0, 1, -1\}$.*

Proof. Take a positive edge (i, j) and a negative edge (k, ℓ), both of which must exist since $\mathcal{H}$ decomposes $\mathcal{P}_G$. Lemma 1.1 and the 4-cycle condition give that $\{i, j\} \cap \{k, \ell\} = \emptyset$ and the two edges are cycle-compatible with some cycle C. Without loss of generality, say $C = (i, j, k, \ell)$. Since $s_{\mathcal{H}}((j, k)) = s_{\mathcal{H}}((\ell, i)) = 0$, one has $a_j = -a_k$ and $a_\ell = -a_i$. So the claim holds true for one pair of edges with opposite sign.

However, we can say a lot more. Take any other positive or negative edge $e = (i', j')$ of G. By using Lemma 1.1 again, we can use a cycle of length 4 between e and either (i, j) or (k, ℓ) (whichever one with the opposite sign as e) in order to show $\{|a_{i'}|, |a_{j'}|\} = \{|a_i|, |a_j|\}$. Consequently, if q is a vertex belonging to either a positive or negative edge, then a_q can take at most two possible absolute values.

Now, since G is connected, any other vertex is connected to one of those vertices via a chain of zero edges. Note that if (k', ℓ') is a zero edge, then $|a_{k'}| = |a_{\ell'}|$. Thus the other vertices cannot introduce any new absolute values, and one has at most two absolute values among all a_i's, as desired.

We now show that we can move the a_i without changing $\mathcal{P}_G^{(+)}$ and $\mathcal{P}_G^{(-)}$ such that there is at most one nonzero absolute value. If not, without loss of generality only a and b exist as absolute values and $a < b$. Note that moving all a_i with absolute value a to 0 does not change any of the signs of the edges and thus the decomposition. So we can assume that besides 0's, there is only one absolute value b. We can then scale all these vertices with values $\pm b$ to ± 1 without changing the signs. $\qquad\square$

Proposition 1.10 and Lemma 1.4 allow us to restrict attention to hyperplanes of the following form:

(1.1) $$\mathcal{H} : \sum_i \pm x_i = 0.$$

An intuitive restatement is that we can consider a hyperplane as an assignment of a *zero*, *negative*, or *positive charge* to each vertex, corresponding to assigning the value of $0, -1,$ or 1 respectively to the weight a_i. We can then read the signs of the edges off the graph by just adding the charges at the two incident vertices.

Example 1.11. We revisit Example 1.9. We can represent the hyperplane $-x_1 + x_4 - x_5 + x_6 = 0$ by assigning positive charges to $\{4, 6\}$ and negative charges to $\{1, 5\}$ and zero charges to the other vertices. This makes $(1, 2)$ a negative edge and $(3, 4)$ a positive edge; the other edges have zero charge.
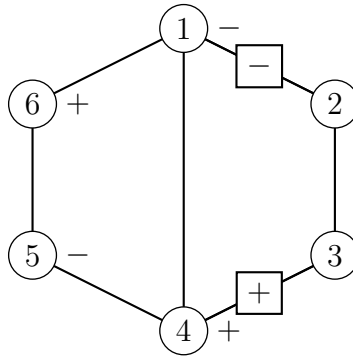
Corollary 1.12. *Suppose we have a decomposition of the edge polytope $\mathcal{P}_G$. Then we can assume one of the following two cases for the vertices of G:*

(1) *There are no vertices with weight* 0. *All positive edges have signature* $\{1, 1\}$ *and all negative edges have signature* $\{-1, -1\}$.

(2) *There is at least one vertex with weight* 0. *All positive edges have signature* $\{1, 0\}$ *and all negative edges have signature* $\{-1, 0\}$.

Proof. If there are no vertices with weight 0, it is easy to see that we are in the first case. It suffices to prove that when we have at least one vertex with weight 0, the conditions in the second case apply. Since G is connected, there must be at least one edge of signature $\{0, a\}$ where $a \in \{-1, 1\}$. Proposition 1.10 gives that all positive edges must then be of signature $\{0, 1\}$ and all negative edges must have signature $\{0, -1\}$. Thus, no edges of signatures $\{1, 1\}$ or $\{-1, -1\}$ can occur, which is equivalent to saying that all nonzero signatures are $\{1, 0\}$ or $\{-1, 0\}$ (equivalently, the 1-weighted vertices and the (-1)-weighted vertices are isolated sets). $\qquad\square$

We call the two types of labelings from Corollary 1.12 *type I* and *type II* respectively. Using these two types, we give an algorithm to check the decomposability of an edge polytope P_G from its graph:

Algorithm 1.13. We check for type *I* and type *II* decomposability separately.

(*type I*) Create an empty list F. For every pair of vertex-disjoint edges e_- and e_+:

(1) set the signatures of e_- and e_+ to $\{-1, -1\}$ and $\{1, 1\}$ respectively by setting the weights of relevant vertices.

(2) try to set weights -1 and 1 to the other vertices one at a time, each time setting the weight of a vertex adjacent to a vertex with weight already set, until one of the following occurs:

  (a) the weights of all vertices are set, in which case we have a decomposition;

  (b) If we have forced any edge e to be a positive or negative edge, check it against every edge f with opposite sign. We cannot continue if one of the following two things happen:

    (i) if e and f are not cycle-compatible (4-cycle condition)

    (ii) if (e, f) is in F (we have failed to assign opposite signs to these edges in the past).

    if one of these happens, add (e_-, e_+) to F and stop the search.

(*type II*) Create an empty list F. For every pair of vertex-disjoint edges e_- and e_+:

(1) set the signatures of e_- and e_+ to $\{-1, 0\}$ and $\{1, 0\}$ respectively by setting the weights of relevant vertices (we do 4 for-loops in this case, corresponding to the 4 possible assignments of the weights).

(2) try to set weights -1, 0, or 1 to the other vertices one at a time, each time setting the weight of a vertex adjacent to a vertex already set with weight -1 or 1 (note this means we always have at most 2 choices, since we cannot put two 1's adjacent to each other or two -1's adjacent to each other by Corollary 1.12), until one of the following occurs:

7

(a) all vertices are set, in which case we have a decomposition;

(b) there are no non-set vertices adjacent to vertices set with -1 and 1, in which case we may set all the unset vertices to weight 0 and obtain a decomposition;

(c) we prune in a similar manner to the type I case (for every new weighted edge, check if it is compatible with the existing edges of the other parity via checking both the 4-cycle condition and F). If we cannot continue, add (e_-, e_+) to F and stop the search.

Naively, Proposition 1.10 tells us we can check decomposability by iterating over 3^n cases. Algorithm 1.13 cuts the base of the exponent to 2 in each of the two types. Though the result is still exponential in the worst possible cases (namely very dense graphs), the pruning process should usually provide large optimizations. Note that checking the compatibility conditions of edges with opposite parity is not a bottleneck of the calculation, as in both types we can keep a precomputed hash for all cycle-compatible edge pairs with a one-time $O(n^4)$ calculation.

Question 1.14. Is there a polynomial-time algorithm to decide if $\mathcal{P}_G$ is decomposable? Is there a polynomial-time algorithm to decide if $\mathcal{P}_G$ is type I (or type II) decomposable?

2. COUNTING DECOMPOSITIONS

Given a graph G, we may also be interested in counting the number of decompositions. Since even the decidability of the existence of a decomposition seems to be difficult, we do not expect this to be a tractable problem except in specific cases; moreover, we need to be careful since different separating hyperplanes may give the same decomposition.

Example 2.1. Consider the four-cycle C_4 with vertices $1, 2, 3, 4$. Then the hyperplanes $x_1 - x_4 = 0$, $x_2 - x_3 = 0$ and $x_1 + x_2 - x_3 - x_4 = 0$ give us the same decomposition of $\mathcal{P}_{C_4}$. Notice that the first two hyperplanes are of type II (with at least one zero coefficient) and the third hyperplane is of type I (with no zero coefficient).

If we restrict to type I hyperplanes, however, we get a unique decomposition (up to sign):

Lemma 2.2. *Two different type I hyperplanes will result in different decompositions of the edge polytope, unless they differ only by an overall sign on the weights.*

Proof. Suppose two hyperplanes give the same decomposition $\mathcal{P}_G = \mathcal{P}_{G_+} \cup \mathcal{P}_{G_-}$. Then up to sign, we can suppose that the edges in $G_0 = G_+ \cap G_-$ are all zero, the edges in $G_+ \backslash G_0$ are all positive, and the edges in $G_- \backslash G_0$ are all negative. Since both hyperplances are of type I, we have $a_i = a_j = 1$ for the coefficients in the

hyperplane for each positive edge (i, j) and $a_k = a_\ell = -1$ for each negative edge (k, ℓ). Then all the other coefficients are uniquely determined. $\square$

Using this fact, we now prove a couple of auxiliary results that combine to count the number of decompositions for complete multipartite graphs.

Lemma 2.3. *In a type II decomposition, no vertex with weight* 1 *may share a neighbor with a vertex of weight* -1.

Proof. By Corollary 1.12, such a neighbor must be 0. However, we now have a positive edge next to a negative one, a contradiction. $\square$

Proposition 2.4. *For a connected bipartite graph G with a type II separating hyperplane, there exists a type I separating hyperplane which gives the same decomposition.*

Proof. Let the bipartition of the vertices be $L \cup R$. By Corollary 1.12, positive edges must have signature $\{1, 0\}$ and negative edges must have signature $\{0, -1\}$. Without loss of generality there is one positive edge (i, j) with $i \in L$ and $j \in R$, such that $a_i = 1$ and $a_j = 0$. Then by Corollary 1.5, all negative edges (k, ℓ) with $a_k = 0$ and $a_\ell = -1$ must have $k \in L$ and $\ell \in R$. By the same reasoning all positive edges must have their weight-1 vertex in L and weight-0 vertex in R.

At this point, we have covered all vertices involved in nonzero edges, so only zero edges remain. Consider a vertex i with weight 1 that we have not yet considered. If it exists, since G is connected, there must exist a path starting at i with weights $1, -1, 1, -1, \ldots$ until we get to a vertex next to a nonzero edge. It is then clear that i must be in L. Using this and a symmetric argument repeatedly, we see that all vertices with weight 1 must be in L and all vertices with weight -1 must be in R.

Finally, consider the following transformation: change all zero weights in R to 1 and all zero weights in L to -1. We claim that no edge will change sign: any positive edge (with signature $\{1, 0\}$) will change to signature $\{1, 1\}$, any negative edge (with signature $\{0, -1\}$) will change to signature $\{-1, -1\}$, any zero edge with signature $\{0, 0\}$ will change to $\{-1, 1\}$, and any zero edge with signature $\{1, -1\}$ is unchanged. The result is then a type II hyperplane which gives the same decomposition. $\square$

Theorem 2.5. *Let G be a complete multipartite graph on vertices $[d] = V_1 \cup \cdots \cup V_k$ with $k \geq 2$. Then the number of decompositions is*

$$2^{d-1} - \sum_i (2^{|V_i|} - 1) - 1.$$

Proof. First, we show that we only need to consider type I hyperplanes. The proof is different for the cases $k = 2$ and $k > 2$: when $k = 2$, we appeal to Proposition 2.4. When $k > 2$, every two vertices share a neighbor, so we use Lemma 2.3.

There are 2^d ways of assigning 1 or -1 to all the vertices. Each of these assignments gives a decomposition unless:

9

(1) all vertices have the same weight, or

(2) both weights occur, but all vertices of one of the weights are inside a single V_i (in which case we do not have an edge of the corresponding sign).

There are 2 assignments for the first case, and $2(2^{|V_i|} - 1)$ ways of confining all weights of some sign (and having at least one) into V_i. Summing and dividing by 2 (using Lemma 2.2) gives the desired answer.

$\square$

In the special case of a complete graph (i.e. $|V_i| = 1$ for all i), this reduces to $2^{d-1} - d - 1$. We remark that this recovers a special case of Theorem 5.3 in [2] when $k = 2$, which studies the splitting of hypersimplices $\Delta(k, n)$.

3. Normal edge polytopes

Recall that an integral convex polytope $\mathcal{P} \subset \mathbb{R}^d$ is *normal* if, for all positive integers N and for all $\beta \in N\mathcal{P} \cap \mathbb{Z}^d$, where $N\mathcal{P} = \{N\alpha : \alpha \in \mathcal{P}\}$, there exists $\beta_1, \ldots, \beta_N$ belonging to $\mathcal{P} \cap \mathbb{Z}^d$ such that $\beta = \sum_i \beta_i$.

Now, consider the following condition on G: If C and C' are cycles of G of odd length, then either they share a vertex or there is an edge (i, j) of G such that $i \in C$ and $j \in C'$ (such an edge is called a *bridge* between C and C'). This condition is called the *odd cycle condition*, which was first investigated by [1] in classical combinatorics. Its relevance to our situation is the following result, which changes the algebraic condition into a combinatorial condition:

Proposition 3.1 (Corollary 2.3 of [4]). $\mathcal{P}_G$ *is normal if and only if G satisfies the odd cycle condition.*

With this result, we now give the main theorem of this section, showing that normality is a hereditary condition on edge polytopes under decomposition:

Theorem 3.2. *Let G be a finite connected simple graph on $[d]$ and suppose that the edge polytope decomposes as $\mathcal{P}_G = \mathcal{P}_{G_+} \cup \mathcal{P}_{G_-}$. Then $\mathcal{P}_G$ is normal if and only if both $\mathcal{P}_{G_+}$ and $\mathcal{P}_{G_-}$ are normal.*

Proof. Recall that G_+ has all the zero and positive edges and G_- has all the zero and negative edges.

Suppose that both $\mathcal{P}_{G_+}$ and $\mathcal{P}_{G_-}$ are normal but $\mathcal{P}_G$ is not. It follows from the odd cycle condition that G contains two disjoint odd cycles C and C' with no bridge. Without loss of generality, we may assume that these are minimal in the sense that we cannot pick a smaller odd cycle to replace C while keeping C' fixed such that they are still disjoint and bridgeless (and vice-versa). In particular, it follows that neither cycle contains a chord, since then a smaller odd cycle could have been selected. If either G_+ or G_- contains all the edges in both cycles, then $\mathcal{P}_G$ cannot be normal. Thus there is at least one positive and at least one negative edge in the edges of C

and C'. If they are in the same cycle C, then since the two edges induce a cycle of length 4, C must have a chord, a contradiction. If they are in different cycles, then a cycle of length 4 between them introduces two bridges that are zero edges, which is also a contradiction. Thus $\mathcal{P}_G$ is normal.

Conversely, suppose that $\mathcal{P}_G$ is normal and that $\mathcal{P}_{G_+}$ is nonnormal. Since $\mathcal{P}_G$ is normal, there is at least one bridge between C and C' in G. Again by the odd cycle condition, the subgraph G_+ contains two disjoint odd cycles C and C' such that no bridge between C and C' belongs to G_+. Therefore, all of the bridges between C and C' in G must be negative edges. We claim all of the edges of C and C' must be zero edges. In fact, if we had an edge e of either C or C' which is nonzero, then since e belongs to G_+ it must be positive, and a cycle of length 4 arises from e and a (negative) bridge, yielding a zero edge bridge between C and C', a contradiction.

Now, let $a_1 x_1 + \cdots + a_d x_d = 0$ define the separating hyperplane $\mathcal{H}$. Since all of the edges of C and C' must be zero edges, one has $a_i + a_j = 0$ if (i, j) is an edge of either C or C'. However, since both C and C' are odd, it follows that $a_i = 0$ for all vertices i of C and C'. Hence there exists no negative bridge between C and C', which gives a contradiction. Hence $\mathcal{P}_{G_+}$ is normal. Similarly, $\mathcal{P}_{G_-}$ is normal. $\square$

Example 3.3. When $\mathcal{P}_G$ is not normal, different types of decompositions can occur, as shown in the following examples:

(1) It is possible for a nonnormal edge polytope to be decomposed into two nonnormal polytopes. Let G be the finite graph with the following decomposition:
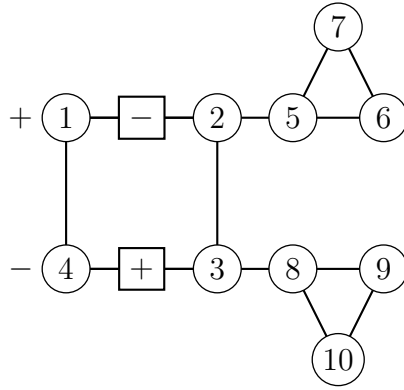
$\mathcal{P}_G$ is not normal because of the two 3-cycles $C_1 = (5, 6, 7)$ and $C_2 = (8, 9, 10)$. Both subpolytopes $\mathcal{P}_{G_+}$ and $\mathcal{P}_{G_-}$ are also nonnormal.

(2) It is possible for a nonnormal edge polytope to be decomposed into one normal polytope and one nonnormal polytope. Let G be the finite graph with the following decomposition:
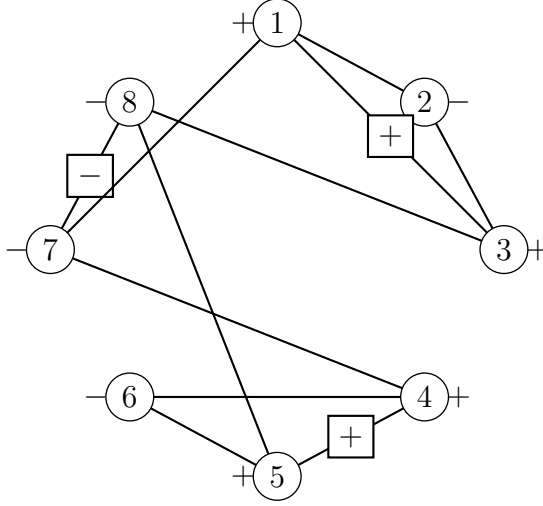
$\mathcal{P}_{G_-}$ is normal and $\mathcal{P}_{G_+}$ is nonnormal.

4. QUADRATIC TORIC IDEALS

Fix a graph G. Let $K[\mathbf{t}] = K[t_1, \ldots, t_d]$ and $K[\mathbf{x}] = K[x_1, \ldots, x_n]$ be the polynomial rings over a field K in d and n variables respectively. For each edge $e = \{i, j\} \in E(G)$, we write $\mathbf{t}^e$ for the squarefree quadratic monomial $t_i t_j \in K[\mathbf{t}]$, and define the *edge ring* $K[G]$ of G to be the subalgebra of $K[\mathbf{t}]$ generated by $\mathbf{t}^{e_1}, \mathbf{t}^{e_2}, \ldots, \mathbf{t}^{e_n}$ over K. Now define the surjective homomorphism of semigroup rings $\pi : K[\mathbf{x}] \to K[G]$ by $\pi(x_i) = \mathbf{t}^{e_i}$ for all $1 \leq i \leq n$. The kernel of π is called the *toric ideal* of G and denoted by I_G. We call a graph G (or its edge polytope $\mathcal{P}_G$) *quadratic* if its toric ideal I_G is generated by quadratic binomials. In the previous section, we studied how the normality property behaves under decomposition; in this section, we do the same with the quadratic property.

Given an even cycle $C = (1, 2, \ldots, k)$:

(1) call (i, j) an *even (odd) chord* of C if $(j - i) \neq 0$ is odd (even). The naming convention corresponds to the parities of the two subcycles created by (i, j) inside C. For example, if $j - i = 2$, then we create a triangle, which is odd. The fact that C is even makes this notion well-defined, as the two subcycles will have the same parity.

(2) call C *long* if it has length at least 6.

(3) given two chords e_1 and e_2 of C, say that they *cross* if they contain 4 distinct vertices that, when we place them in order around C, gives the sequence (a_1, a_2, a_3, a_4) where the set of edges $\{e_1, e_2\}$ is the set $\{(a_1, a_3), (a_2, a_4)\}$.

(4) call a triple of chords $S = (c_1, c_2, c_3)$ of C an *odd-triple* if at least two of them *cross*.

12

As before, our strategy is to convert the criterion of being quadratic to that of a combinatorial criterion (Theorem 4.1) on graphs, which we will use to prove our final main result, Theorem 4.3.

Theorem 4.1 (Theorem 1.2 of [7])**.** *G is quadratic if and only if the following two conditions hold:*

(1) *For every even long cycle, there exists either an even chord or an odd-triple.*
(2) *The induced graph of any two odd-cycles are 2-connected by bridges (namely, for any two odd cycles which share exactly one node, there exists a bridge which does not go through the common node, and for any two odd cycles which do not share any nodes, there exist two bridges).*

Lemma 4.2. *Let C be an even cycle inside a quadratic G, with edges $e_1, e_2, \ldots, e_{2k}$. Suppose C contains exactly one positive e_i and one negative edge e_j. Then $j \cong i$ (mod 2).*

Proof. Suppose not. Without loss of generality e_1 is positive and e_{2l} is negative. We can assume that C has vertices $c_1, \ldots, c_{2k}$ with $e_i = (c_i, c_{i+1})$, where it is understood the vertices are labeled modulo $2k$, and that either c_1 or c_2 is 1. Without loss of generality $c_2 = 1$. Since the edges $e_2, \ldots, e_{2l-1}$ are all zero edges, we must have $c_3 = -1, c_4 = 1, c_5 = -1, \ldots, c_{2l} = 1$, which is a contradiction because $e_{2l} = (c_{2l}, c_{2l+1})$ is negative. $\square$

Theorem 4.3. *Let $\mathcal{P}_G = \mathcal{P}_{G_+} \cup \mathcal{P}_{G_-}$ be a decomposition. If $\mathcal{P}_{G_+}$ and $\mathcal{P}_{G_-}$ are both quadratic, then $\mathcal{P}_G$ is quadratic.*

Proof. Suppose $\mathcal{P}_G$ were not quadratic. Then G fails at least one property listed in Theorem 4.1. This means one of the following must be true:

(1) G has a even long cycle C without even chords or odd triples.
(2) G has two edge-disjoint odd cycles C_1 and C_2 that have at most one bridge between them and do not share a vertex.
(3) G has two edge-disjoint odd cycles C_1 and C_2 that have no bridges and share a vertex.

We'll show that each of these cases causes a contradiction:

(1) There must be at least one positive edge e_1 in C, since otherwise C would appear in G_-, failing the assumption that G_- was quadratic. Similarly, C must contain a negative edge e_2 as well. There must be a 4-cycle (e_1, u, e_2, v), so u and v are a pair of chords of C with the same parity. By assumption on C having no even chords, u and v must be odd chords of C. Thus, by Lemma 4.2, there must be at least one more nonzero edge e_3 in C. Without loss of generality it is negative; then e_3 forms a 4-cycle with e_1 as well and creates another pair of odd chords C by the same reasoning above and one

13

of them must intersect one of u and v. However, this creates a odd triple inside C, which is a contradiction.
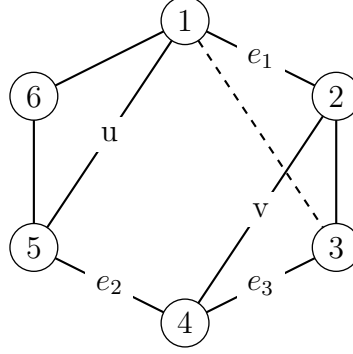

(2) Suppose d is a chord in C_1; then d forms a smaller odd cycle with a subset of vertices of C_1 that must still have at most one bridge with C_2 and share no vertices. Thus, we can assume both C_1 and C_2 are chordless.

   By the same reasoning as the previous case of the long even cycle, there must be at least one positive edge e_1 (and negative edge e_2) in C_1 or C_2, else G_- (respectively G_+) would not be quadratic. First, note that e_1 and e_2 must not be in different cycles, since this would create a 4-cycle between them that would serve as 2 bridges, creating a contradiciton. Thus, all the nonzero edges must be in one of the cycles, without loss of generality C_1. But this is impossible, since the 4-cycle created by e_1 and e_2 must introduce two chords in C_1, which we assumed to be chordless. Thus, we again have a contradiction.

(3) Again, if d is a chord in C_1, then either we get a smaller odd cycle that does not share a vertex or any bridges with C_2, or we get a smaller cycle that shares a vertex with C_2 and no bridges. Thus, we can again assume both C_1 and C_2 are chordless. The same reasoning as above gives a contradiction. $\square$

Example 4.4. Knowing now that we cannot have both G_+ and G_- be quadratic while G is not quadratic, we end this paper with examples showing that all other possibilities under decomposition can be realized. This shows that we cannot extend the statement of Theorem 4.3 to include the "only if" case.
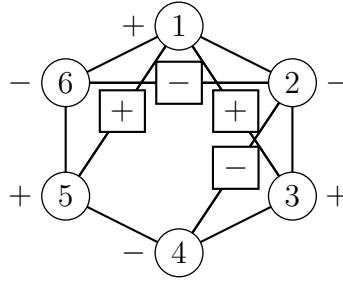
(1) G is quadratic, G_+, G_- are quadratic. Note that this is Example 1.9 again.
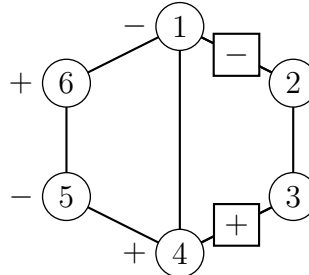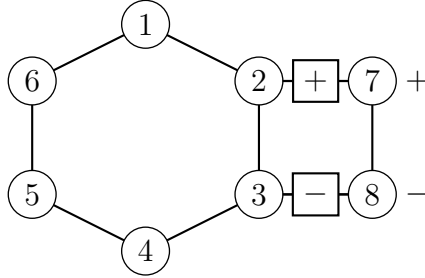

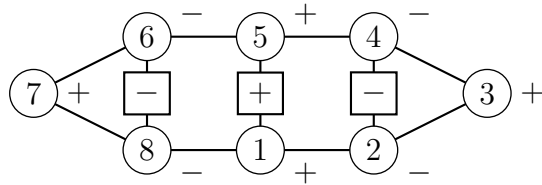(2) G is quadratic, G_+ is not quadratic, and G_- is quadratic.
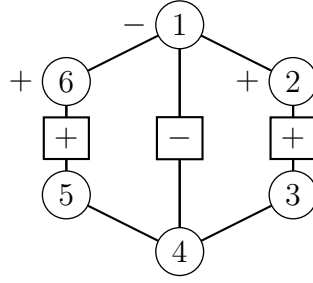
(3) G is quadratic, G_+ and G_- are not quadratic.

(4) G is not quadratic, G_+ and G_- are not quadratic.

(5) G is not quadratic, G_+ is quadratic, and G_- is not quadratic.

References

[1] D. R. Fulkerson, A. J. Hoffman and M. H. McAndrew, Some properties of graphs with multiple edges, *Canad. J. Math.* **17** (1965), 166–177.

[2] S. Herrmann and M. Joswig, Splitting polytopes, *Münster J. of Math.* **1** (2008), 109–142.

[3] M. Hochster, Rings of invariants, Cohen-Macaulay rings generated by monomials, and polytopes, *Ann. Math.* **96** (1972), 318–338.

[4] H. Ohsugi and T. Hibi, Normal Polytopes Arising from Finite Graphs, *J. Algebra* **207** (1998), 409–426.

[5] H. Ohsugi and T. Hibi, A normal $(0,1)$-polytope none of whose regular triangulations is unimodular, *Discrete Comput. Geom.* **21** (1999), 201–204.

[6] H. Ohsugi and T. Hibi, Koszul bipartite graphs, *Advances in Applied Math.* **22** (1999), 25–28.

[7] H. Ohsugi and T. Hibi, Toric ideals generated by quadratic binomials, *J. Algebra* **218** (1999), 509–527.

[8] H. Ohsugi and T. Hibi, Compressed polytopes, initial ideals and complete multipartite graphs, *Illinois J. Math.* **44** (2000), 391–406.

[9] H. Ohsugi and T. Hibi, Simple polytopes arising from finite graphs *ITSLCSREA Press* **44** (2008), 73–79.

[10] A. Simis, W. V. Vasconcelos and R. H. Villarreal, The integral closure of subrings associated to graphs, *J. Algebra* **199** (1998), 281–289.

[11] B. Sturmfels, "Gröbner bases and convex polytopes," Amer. Math. Soc., Providence, RI, 1995.

Takayuki Hibi, Department of Pure and Applied Mathematics, Graduate School of Information Science and Technology, Osaka University, Toyonaka, Osaka 560-0043, Japan

E-mail address: `hibi@math.sci.osaka-u.ac.jp`

Nan Li, Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

E-mail address: `nan@math.mit.edu`

Yan X Zhang, Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

E-mail address: `yanzhang@math.mit.edu`